On the Average Random Probing Model

Julien Béguinot[1] and Loïc Masure[2†]

[1] LTCI, Télécom Paris, Institut Polytechnique de Paris, julien.beguinot@telecom-paris.fr
[2] LIRMM, Univ. Montpellier, CNRS, loic.masure@lirmm.fr

Abstract. Masking is one of the main countermeasures against side-channel analysis since it relies on provable security. In this context, "provable" means that a security bound can be exhibited for the masked implementation through a theoretical analysis in a given threat model. The main goal in this line of research is therefore to provide the tightest security bound, in the most realistic model, in the most generic way. Yet, all of these objectives cannot be reached together. That is why the masking literature has introduced a large spectrum of threat models and reductions between them, depending on the desired trade-off with respect to these three goals. In this paper, we focus on three threat models, namely the noisy-leakage model (realistic yet hard to work with), the random probing (unrealistic yet easy to work with), and more particularly a third intermediate model called average random probing. Average random probing has been introduced by Dziembowski et al. at Eurocrypt 2015, in order to exhibit a tight reduction between noisy-leakage and random probing models, recently proven by Brian *et al.* at Eurocrypt 2024. This milestone has strong practical consequences, since otherwise the reduction from the noisy leakage model to the random probing model introduces a prohibitively high constant factor in the security bound, preventing security evaluators to use it in practice. However, we exhibit a gap between the average random probing definitions of Dziembowski *et al.* (denoted hereafter by DFS-ARP) and Brian *et al.* (simply denoted by ARP). Whereas any noisy leakage can be tightly reduced to DFS-ARP, we show in this paper that it cannot be tightly reduced to ARP, unless requiring extra assumptions, *e.g.*, if the noisy leakage is deterministic. Our proof techniques do not involve more tools than the one used so far in such reductions, namely basic probability facts, and known properties of the total variation distance. As a consequence, the reduction from the noisy leakage to the random probing — without high constant factor — remains unproven. This stresses the need to clarify the practical relevance of analyzing the security of masking in the random probing model since most of the current efforts towards improving the constructions and their security proofs in the random probing model might be hindered by potentially unavoidable loss in the reduction from more realistic but currently less investigated leakage models.

Keywords: Masking · Noisy leakage · Random Probing · Average Random Probing · Reduction · Leakage Model

1 Introduction

Context. In the quest of provably secure implementations against side-channel analysis, the works of Duc, Dziembowski and Faust [DDF14] marked an important milestone, bridging the gap between two different visions of proving the security of the masking countermeasure [GP99].

†The authors would like to thank Gianluca Brian, Stefan Dziembowski, Sebastian Faust, Elena Micheli, Maximilian Orlt, and François-Xavier Standaert for the fruitful discussions on this topic. This work received funding from the France 2030 program, managed by the French National Research Agency under grant agreement No. ANR-22-PETQ-0008 PQ-TLS.

On the one hand, the vision of Ishai, Sahai and Wagner (ISW) relies on the *probing* security of a logical or arithmetical circuit [ISW03]. That is, in a circuit implemented according to ISW, any set of at most $(d-1)$ intermediate computations remains independent of any secret fed as a d-secret sharing to the input of the circuit. This security notion, while depicting an unrealistic adversary able to exactly probe a limited amount of wires in the circuit, has the merits to be easy to verify [BMRT22], interpret [BDF+17], and even compose to some extent [BBD+16]. However, this approach cannot prove the ineffectiveness of all attacks, such as horizontal attacks [BCPZ16]. Even worse, some masking schemes such as table re-computation are even provably secure in the probing model [Cor14], while exposing critical vulnerability in practice [BS21]. Therefore, researchers have looked for another way to characterize the soundness of masking.

On the other hand, the seminal approach triggered by Chari *et al.* [CJRR99], and later more formalized by Prouff and Rivain [PR13], establishes information-theoretic security bounds — thereby making masking a *universal* countermeasure — in a much more realistic leakage model. Here, the leakage is said δ-*noisy*, if each wire independently reveals some leakage to the adversary such that the statistical bias between the wire and its corresponding leakage is bounded by a scalar parameter $\delta \in [0, 1]$.[1] This requires, however, much more tedious proof strategies, requiring strong assumptions on some *leak-free* computations, too unrealistic to faithfully depict the actual security of some masking schemes [CPRR14].

Duc *et al.*'s unification came timely to bridge the gap between these two visions, by means of an intermediate leakage model, called *random probing* (RP). In this model, each wire is revealed to the adversary with probability $\epsilon_{\mathsf{RP}} \in [0, 1]$. Duc *et al.* proved first that the δ-noisy leakage model can be perfectly simulated by the ϵ_{RP}-random probing model, provided that $\epsilon_{\mathsf{RP}} = \delta \cdot |\mathcal{X}|$, where $\mathcal{X}$ is the set of all values that any wire may take in the circuit, *e.g.*, the field size in an arithmetical circuit.[2] Secondly, they proved that the ϵ_{RP}-random probing model can in turn be simulated by the $(d-1)$-probing model with some error bounded by $(\epsilon_{\mathsf{RP}} \cdot \mathcal{O}(d))^d$ in the case of the ISW construction [Man23, Lemma 14]. Overall, this shows that a side-channel adversary has no more advantage than $\xi = (\delta \cdot \mathcal{O}(d) \cdot |\mathcal{X}|)^d$ compared to a blind-guess adversary.

This result justifies that masking a cryptographic implementation is *sound* from a theoretic point of view. However, to the best of our knowledge the security bound ξ in itself has never been used by security evaluators or designers to claim a given security level, since the product $(\mathcal{O}(d) \cdot |\mathcal{X}|)^d$ is prohibitively high. In typical applications, $d = 2, 3, 4$, whereas $|\mathcal{X}| = 2^8$ for the AES, $\approx 2^{12}$ for Kyber, $\approx 2^{23}$ for Dilithium, or even $\approx 2^{49}$ for the masking-friendly lattice-based signature Raccoon. A recent line of research, conducted by Belaïd *et al.* endeavoured to substitute the $\mathcal{O}(d)^d$ factor in the security bound ξ with $\mathcal{O}(1)^d$, by means of an alternative masking construction based on the so-called *expansion* strategy, directly proven secure in the random probing model [BCP+20, BRT21].

However, getting rid of the $|\mathcal{X}|^d$ factor remains the main challenge towards making the security bound practical. Hereupon, much less attempts have been made in the literature. To date, two approaches have been proposed. The first one consists in changing the way in which the noisy-leakage parameter δ is measured, as proposed by Prest *et al.* [PGMP19] and later optimally improved by Béguinot *et al.* [BCGR24]. But this requires using more conservative metrics that potentially hide the reduction loss factor inside the metric [MS23]. Moreover, in practice such metrics are less convenient to compute, since they require to estimate the extrema of the leakage distribution based on empirical data, whereas Duc *et al.*'s noise parameter only requires estimating averages.

The second explored approach consists in using the *Average Random Probing* model,

[1] Such a δ parameter can be efficiently measured by a side-channel security evaluator, when characterizing the target device.

[2] We provide concrete examples in the next paragraph.

introduced by Dziembowski, Faust and Skorski [DFS15b], as a surrogate to the random probing (RP). In this model, each wire is revealed to the adversary with probability ϵ_x, now allowed to depend on the underlying value x carried by the wire. Dziembowski *et al.* have shown that this relaxed version of the RP and the noisy leakage model are essentially equivalent, in the sense that $\epsilon_{\mathsf{DFS}} = \mathop{\mathbb{E}}_{x}[\epsilon_x] \leq \delta$ — hence the "average" terminology. In other words, this avoids the reduction loss induced by the direct reduction from the noisy leakage to the RP. The average random probing model did not get much attention since then, until Brian *et al.* recently exhibited a reduction from this model to the RP, thereby tightening the chain of reductions from the noisy leakage (NL) to the probing model [BDF24]. Concretely, it would allow to remove the $|\mathcal{X}|^d$ factor, at the cost of doubling the number of shares, by transforming any random-probing secure circuit according to a circuit compiler designed by Brian *et al.* This ground-breaking result underlines the relevance of the average random probing, as a corner-stone of masking security proofs.

The Issue. The starting point of this paper is however to point out at a slight gap between the definitions of the average random probing provided by Dziembowski *et al.* [DFS15b] on the one hand (denoted hereafter by DFS-ARP), and the average random probing model defined by Brian *et al.* [BDF24] (simply denoted by ARP hereafter) on the other hand.[3] In a nutshell, Dziembowski *et al.* tweaked the ARP oracle revealing each wire to the adversary with probability ϵ_x, so that it also reveals some internal randomness used to stochastically decide whether each wire must be revealed or not. Surprisingly, this internal randomness turns out to carry a lot of information on the secret value of the wire, which would not be provided to an ARP adversary. This observation, counter-intuitive at first glance, results from the fact that this internal randomness, once drawn offline, is compared to a sensitive value during a rejection sampling before being passed to the simulator. This creates a bias between the internal randomness returned along with a revealed wire and that returned alone. Unfortunately, this bias turns out to depend on the underlying wire value — we elaborate more on that in Subsection 4.2.

Concretely, this extra information is so critical that the simulator used by Dziembowski *et al.* to prove the tight reduction from NL to DFS-ARP no longer applies when considering the ARP model nowadays. This initial observation therefore triggers a natural question:

> *Is it still possible to simulate any noisy leakage from some ARP leakage? And if so, at what cost in the reduction?*

Our Contribution. We address this question with a three-fold answer: (1) contrary to the RP, any non-injective leakage function is indeed reducible to the ARP in a non-trivial manner; (2) the reduction is not tight generally speaking, as we exhibit a counterexample where the reduction to ARP requires $\epsilon_{\mathsf{ARP}} = \Omega(\delta \cdot |\mathcal{X}|)$; (3) still, for particular cases such as deterministic leakage models, we prove that the reduction is tight, whereas such noisy leakage models cannot even be simulated in the random probing model.

To obtain our results, we revisit the proof of the core technical result of Duc *et al.* [DDF14], namely the lemma reducing the noisy leakage to the random probing, in light of the relaxed definition of ARP. We derive a new necessary and sufficient condition for simulation in the ARP. This grounds the stage for exhibiting the so-called *catastrophic* channel as an exemplary leakage function to disprove the tightness between the noisy leakage and the ARP, and for proving the tight reduction in two particular use-cases. Our proofs do not involve more tools and techniques than the one used by Duc *et al.* [DDF14], namely basic probability facts, and functional properties of the statistical bias.

[3] We take the convention to consider the definition of Brian *et al.* as the most natural.

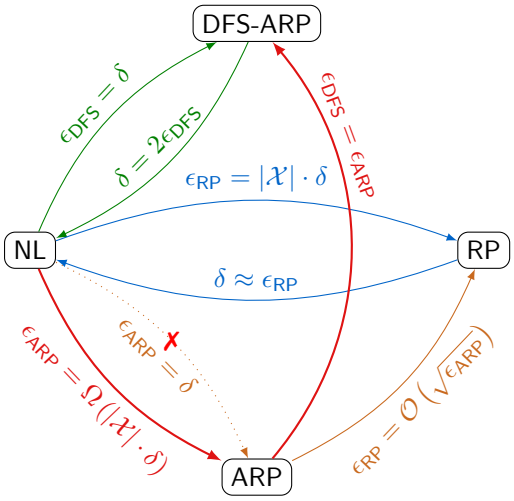

Figure 1: The leakage models, and their relationships established in the literature. Blue: [DDF14]. Green: [DFS15b]. Brown: [BDF24]. The dotted curve corresponds to the incorrect claim of [BDF24, Thm. 3]. Red: our contribution.

Impact & Perspectives. Even though our works do not disprove any of the theoretical results of Dziembowski *et al.* [DFS15b] nor the main theorem of Brian *et al.* [BDF24], they invalidate the essential interest of their combination, as stated in [BDF24, Cor. 9]. In other words, the tight reduction from the noisy leakage to the random probing remains unproven, except in the particular cases where we exhibited a tight reduction from the noisy leakage model to the ARP — see Section 6. We summarize the whole picture in Figure 1.

Overall, this denotes the need to clarify the practical relevance of analyzing the security of masking in the random probing model [DFZ19, BCP+20, CFOS21, BRT21, BRTV21, BFO23, JMB24, WJSW24]. The impossible tightness between the noisy leakage model and the ARP does not necessarily invalidate the hope for a tight reduction to the random probing model. We even believe that some ideas in Brian *et al.*'s techniques [BDF24] could be leveraged to this end, and we elaborate more on that in Section 7. Yet, this suggests at least that the side-channel and provably-secure masking research community should pay more attention to this current *elephant in the room*, since most of the current efforts towards improving the constructions and their security proofs in the random probing model might be hindered by potentially unavoidable loss in the reduction from more realistic — but currently less investigated — leakage models.

Organization of the Paper. We first provide some background and recall some definitions of the different leakage models in Section 2 and Section 3. We then introduce the (different flavors of the) average random probing model in Section 4, and emphasize the issue of the model as defined by Dziembowski *et al.* We then make a thorough characterization of the average random probing model, as defined by Brian *et al.* in Section 5, in order to derive our negative results. Then we exhibit in Section 6 two use-cases where the reduction from noisy leakage to ARP remains tight. We conclude on some discussions and perspectives in Section 7.

2 Preliminaries

Let X, Y be two random variables. We note $\mathrm{X} \overset{d}{=} \mathrm{Y}$ if their probability distributions are equal, *i.e.*, $\Pr(\mathrm{X}) = \Pr(\mathrm{Y})$. In this paper, we manipulate the *total variation distance*, which measures the dissimilarity between two probability mass functions (pmfs).

Definition 1 (Total Variation Distance)**.** Let p, m be two pmfs over a finite set $\mathcal{L}$. The *total variation* distance between p and m, denoted by $\mathsf{TV}(\mathsf{p}; \mathsf{m})$, is defined as follows:

$$\mathsf{TV}(\mathsf{p}; \mathsf{m}) = \frac{1}{2} \sum_{l \in \mathcal{L}} |\mathsf{p}(l) - \mathsf{m}(l)| \quad .$$

In the literature, the total variation distance is also known as the *statistical distance*. Likewise, we define the statistical *bias* induced by some random variable on another.

Definition 2 (Statistical Bias)**.** The *statistical bias* between two random variables X, L is defined by the following equality:

$$\Delta(\mathrm{X}; \mathrm{L}) = \mathsf{TV}(\Pr(\mathrm{L}, \mathrm{X}); \Pr(\mathrm{L}) \otimes \Pr(\mathrm{X})),$$

where $\otimes$ denotes the Cartesian product: $(\Pr(\mathrm{L}) \otimes \Pr(\mathrm{X}))(l, x) = \Pr(\mathrm{L} = l) \cdot \Pr(\mathrm{X} = x)$.

The statistical bias is used in this paper to characterize the amount of leakage from each wire of the circuit, as formalized hereafter.

Definition 3 (Noisy Leakage)**.** A (possibly randomized) leakage function $\mathrm{L} : \mathcal{X} \to \mathcal{L}$ is said to be δ-*noisy* if $\Delta(\mathrm{X}; \mathrm{L}) \leq \delta$.

We provide hereafter a few other identities verified by the metrics.

Proposition 1. *The total variation and the statistical bias may be equivalently be expressed as follows:*[4]

$$\mathsf{TV}(\mathsf{p}; \mathsf{m}) = \sum_{l \in \mathcal{L}} \max\{0, \mathsf{p}(l) - \mathsf{m}(l)\} = 1 - \sum_{l \in \mathcal{L}} \min\{\mathsf{p}(l), \mathsf{m}(l)\} = \max_{T \subseteq \mathcal{L}} |\mathsf{p}(T) - \mathsf{m}(T)| \quad ,$$

$$\Delta(\mathrm{X}; \mathrm{L}) = \mathop{\mathbb{E}}_{x} [\mathsf{TV}(\Pr(\mathrm{L} \mid \mathrm{X} = x); \Pr(\mathrm{L}))] = \mathop{\mathbb{E}}_{l} [\mathsf{TV}(\Pr(\mathrm{X} \mid \mathrm{L} = l); \Pr(\mathrm{X}))] \quad .$$

Data-processing inequality*: Let $\mathrm{L}', \mathcal{S}$ be two functions such that $\mathcal{X} \overset{\mathrm{L}'}{\to} \mathcal{Y} \overset{\mathcal{S}}{\to} \mathcal{L}$:* $\Delta(\mathrm{X}; \mathcal{S} \circ \mathrm{L}'(\mathrm{X})) \leq \Delta(\mathrm{X}; \mathrm{L}'(\mathrm{X}))$.

3 Warm-Up: the Random Probing Model

Before diving into the details of the average random probing, let us recall the reduction from the noisy leakage model to the random probing model. Later in this paper, we rely on some parts of the proof of the reduction to discuss the reduction to the average random probing, as they share many points.

We start with the definition of the simulation in the random probing model.

Definition 4 (Random Probing)**.** A noisy function $\mathrm{L} : \mathcal{X} \to \mathcal{L}$ is said to be *simulatable in the ϵ-random probing model* if there exists a randomized function $\mathcal{S}$ — the *simulator* — such that for every $x \in \mathcal{X}$, we have

$$\mathrm{L}(x) \overset{d}{=} \mathcal{S}(\varphi(x)) \quad ,$$

[4] The interested reader may refer to Reyzin's lecture notes.

where $\varphi : \mathcal{X} \to \mathcal{X} \cup \{\perp\}$ is the ϵ-identity function, *i.e.*:

$$\varphi(x) = \begin{cases} x, & \text{with probability } \epsilon \in [0,1], \\ \perp, & \text{otherwise} \end{cases} .$$

We say that L is *RP-simulatable* if there exists some $\epsilon \in [0,1)$ for which L is simulatable in the ϵ-random probing.[5]

Next, we recall the key result of Duc, Dziembowski and Faust. The statement of Theorem 1 slightly differs from the one of Duc *et al.*, as we emphasize that the constraint on ϵ is the best that can be obtained.

Theorem 1 ([DDF14, Lemma 2, restated])**.** *Let* $\mathrm{L} : \mathcal{X} \to \mathcal{L}$ *be a (possibly randomized) leakage function. Then,* L *is simulatable in the* ϵ*-random probing model if* and only if

$$\epsilon \geq 1 - \sum_{l} \min_{x} \Pr\left(\mathrm{L}(x) = l\right) .$$

For self-completeness, we revisit the proof of Theorem 1 in order to show how the construction of the simulator naturally emerges from its constraints.

Proof. We proceed by analysis-synthesis. Assume first that such a simulator $\mathcal{S}$ exists, we need to specify for all $x \in \mathcal{X}$ and for all $l \in \mathcal{L}$ the probabilities $\Pr\left(\mathcal{S}\left(x\right) = l\right)$, and $\Pr\left(\mathcal{S}\left(\perp\right) = l\right)$, verifying the following constraints:

1. For all input $x \in \mathcal{X}$, the mapping $l \mapsto \Pr\left(\mathcal{S}\left(x\right) = l\right)$ should be a pmf, *i.e.*,
   (a) For all $l \in \mathcal{L}$, $0 \leq \Pr\left(\mathcal{S}\left(x\right) = l\right)$
   (b) $\sum_l \Pr\left(\mathcal{S}\left(x\right) = l\right) = 1$.
2. For the input $\perp$, the mapping $l \mapsto \Pr\left(\mathcal{S}\left(\perp\right) = l\right)$ should be a pmf, *i.e.*
   (a) For all $l \in \mathcal{L}$, $0 \leq \Pr\left(\mathcal{S}\left(\perp\right) = l\right)$
   (b) $\sum_l \Pr\left(\mathcal{S}\left(\perp\right) = l\right) = 1$.
   (c) The pmf $\Pr\left(\mathcal{S}\left(\perp\right)\right)$ should not depend on any x.
3. For any x, l, $\Pr\left(\mathcal{S}\left(\varphi(x)\right) = l\right) = \Pr\left(\mathrm{L}(x) = l\right)$

Let us start from the last constraint item 3. By using the total probability formula, and by definition of the ϵ-identity function, we have that for any x and any l:

$$\begin{aligned} \Pr\left(\mathrm{L}(x) = l\right) &= \Pr\left(\mathcal{S}\left(\varphi(x)\right) = l\right) \\ &= \Pr\left(\varphi(x) = x\right) \cdot \Pr\left(\mathcal{S}\left(x\right) = l\right) + \Pr\left(\varphi(x) = \perp\right) \cdot \Pr\left(\mathcal{S}\left(\perp\right) = l\right) \\ &= \epsilon \cdot \Pr\left(\mathcal{S}\left(x\right) = l\right) + (1 - \epsilon) \cdot \Pr\left(\mathcal{S}\left(\perp\right) = l\right) . \end{aligned}$$

If $\epsilon = 1$, then the simulator is trivial: $\mathcal{S}\left(x\right)$ samples l according to $\Pr\left(\mathrm{L}(x)\right)$, and we do not need to specify $\Pr\left(\mathcal{S}\left(\perp\right)\right)$, as "$\perp$" is never received by the simulator. Now assume for the rest of the proof that $\epsilon < 1$. We may rephrase the pmf of the simulator, upon receiving "$\perp$", as follows:

$$\Pr\left(\mathcal{S}\left(\perp\right) = l\right) = \frac{\Pr\left(\mathrm{L}(x) = l\right) - \epsilon \cdot \Pr\left(\mathcal{S}\left(x\right) = l\right)}{1 - \epsilon} \tag{1}$$

Note that according to item 2c, the right-hand side of Equation 1 should not depend on any x. Since $1 - \epsilon$ is also independent of x, it implies that the quantity

$$\pi(l) = \Pr\left(\mathrm{L}(x) = l\right) - \epsilon \cdot \Pr\left(\mathcal{S}\left(x\right) = l\right) \tag{2}$$

[5] We exclude the case $\epsilon = 1$ as the simulation would be trivial.

should be independent of x. Moreover, by summing Equation 2 over $\mathcal{L}$, it turns out that $\sum_{l\in\mathcal{L}} \pi(l) = 1 - \epsilon$. Hence, assuming $\epsilon < 1$ — the simulation is trivial otherwise — we may rephrase every pmf of the simulator as functions of the pmf $\Pr(\mathrm{L}(x))$ — which is known — and the quantities $\pi(l)$:

$$\Pr(\mathcal{S}(\perp) = l) = \frac{\pi(l)}{\sum_{l'\in\mathcal{L}} \pi(l')} \ , \tag{3}$$

$$\Pr(\mathcal{S}(x) = l) = \frac{\Pr(\mathrm{L}(x) = l) - \pi(l)}{1 - \sum_{l'\in\mathcal{L}} \pi(l')} \ . \tag{4}$$

We now need to find constraints on $\pi(l)$. To verify item 2a, one must have $\Pr(\mathcal{S}(\perp) = l) \geq 0$, *i.e.*, all the $\pi(l)$ must have the same sign. Given that $\sum_{l\in\mathcal{L}} \pi(l) = 1 - \epsilon \in]0, 1]$, we get the first constraint that

$$0 \leq \pi(l).$$

Moreover, to fulfill item 1a, one must have $\Pr(\mathcal{S}(x) = l) \geq 0$ for all input x, *i.e.*, $\pi(l) \leq \Pr(\mathrm{L}(x) = l)$. Since $\pi(l)$ is independent of x, this implies that

$$\pi(l) \leq \min_{x\in\mathcal{X}} \Pr(\mathrm{L}(x) = l) \, .$$

It therefore remains to verify item 1b and item 2b. item 2b is trivially verified by virtue of Equation 3. As per item 1b, it is verified by virtue of Equation 4, and leveraging the fact that $\sum_l \Pr(\mathrm{L}(x) = l) = 1$ for all input x, as $\Pr(\mathrm{L}(x))$ is a pmf.

As a result, since none of the implications developed so far are contradicting with each other, any randomized function $\mathcal{S}$ defined by Equation 3 and Equation 4, where $0 \leq \pi(l) \leq \min_{x\in\mathcal{X}} \Pr(\mathrm{L}(x) = l)$ for any leakage value l, will result in a perfect simulation of L by $\mathcal{S} \circ \varphi$, where φ is an ϵ-identity and $\epsilon = 1 - \sum_{l\in\mathcal{L}} \pi(l)$. □

Remark 1. From a security analysis point of view, one would like to build a simulator for the smallest ϵ value as possible, denoted thereafter by ϵ_{RP}, *i.e.*, $\epsilon_{\mathsf{RP}} = 1 - \sum_{l\in\mathcal{L}} \min_{x\in\mathcal{X}} \Pr(\mathrm{L}(x) = l)$. This can be done by choosing the greatest possible value for $\pi(l)$, *i.e.*, $\min_{x\in\mathcal{X}} \Pr(\mathrm{L}(x) = l)$.

The following lemma, proven as part of the proof of [DDF14, Lemma 2], links the smallest ϵ_{RP} such that a leakage function is simulatable in the ϵ_{RP}-random probing model, to the statistical bias δ and the input set size.

Lemma 1 ([DDF14])**.** *Let* $\mathrm{L} : \mathcal{X} \to \mathcal{L}$*, be a* δ*-noisy function. Then,* $\epsilon_{\mathsf{RP}} \leq \delta \cdot |\mathcal{X}|$ *.*

The dependency on the field size is at the core of the looseness of the reduction from the noisy leakage to the (region) probing model. Unfortunately, Dziembowski, Faust and Skorski have shown that the field-size factor is unavoidable in general [DFS15b, p. 170]. That is why several *ad hoc* surrogate metrics to the statistical bias have been proposed in the literature.

Proposition 2. *Let* $\epsilon_{\mathsf{RP}} = 1 - \sum_{l\in\mathcal{L}} \min_{x\in\mathcal{X}} \Pr(\mathrm{L}(x) = l)$ *be the optimal parameter of the random probing oracle, for a given noisy leakage* L*. Then,*

$$\epsilon = \mathsf{CDC}(\mathrm{X}; \mathrm{L}) \leq \mathsf{ARE}(\mathrm{X}; \mathrm{L}) \ ,$$

where $\mathsf{CDC}(\mathrm{X}; \mathrm{L}) = \mathop{\mathbb{E}}_{l} \left[\max_{x\in\mathcal{X}} \left(1 - \frac{\Pr(\mathrm{X}=x \mid \mathrm{L}=l)}{\Pr(\mathrm{X}=x)}\right)\right]$ *is the* Complementary Doeblin Coefficient *(CDC) [BCGR24, Def. 9], and* $\mathsf{ARE}(\mathrm{X}; \mathrm{L}) = \mathop{\mathbb{E}}_{l} \left[\max_{x\in\mathcal{X}} \left|1 - \frac{\Pr(\mathrm{X}=x \mid \mathrm{L}=l)}{\Pr(\mathrm{X}=x)}\right|\right]$ *is the* Average Relative Error *[PGMP19, Def. 3].*

Proof. Observe that the probability $\Pr(\mathrm{L}(x) = l)$ can be rephrased as the conditional distribution $\Pr(\mathrm{L} = l | \mathrm{X} = x)$. Let $\Pr(\mathrm{X})$ be any marginal distribution over the random variable X. By applying Bayes' theorem, we have[6]

$$\begin{aligned}\varepsilon = 1 - \sum_l \min_x \Pr(\mathrm{L}(x) = l) &= 1 - \sum_l \min_x \Pr(\mathrm{L} = l \mid \mathrm{X} = x) \\ &= 1 - \sum_l \Pr(\mathrm{L} = l) \min_x \frac{\Pr(\mathrm{X} = x \mid \mathrm{L} = l)}{\Pr(\mathrm{X} = x)} \quad .\end{aligned}$$

One recognizes the expectation over the leakage of the quantity $\min_{x \in \mathcal{X}} \frac{\Pr(\mathrm{X}=x \mid \mathrm{L}=l)}{\Pr(\mathrm{X}=x)}$. Including the constant term "1" into the expectation gives the first equality. The second equality holds by definition of the absolute value. □

4 The Average Random Probing

To circumvent the issue of simulating from the random probing model, a surrogate leakage model called *average random probing* (ARP) has been introduced in the literature. The core idea is to relax the definition of the ϵ-identity function, by allowing a different probability of returning "$\perp$" for each input x, provided that on *average*, the ϵ-identity does not return "$\perp$" with probability more than ϵ. The hope is hence to decrease ϵ enough to get rid of the field-size dependency. We give hereafter a more formal definition. In the remaining of this paper, $(\epsilon_x)_{\mathcal{X}}$ denotes a $|\mathcal{X}|$-dimensional vector in $[0,1]^{|\mathcal{X}|}$.

Definition 5 (Average Random Probing)**.** A noisy function $\mathrm{L} : \mathcal{X} \to \mathcal{L}$ is said to be *simulatable in the $(\epsilon_x)_x$-average random probing model* if there exists a randomized function $\mathcal{S} : \mathcal{X} \cup \{\perp\} \to \mathcal{L}$ — the *simulator* — such that for every $x \in \mathcal{X}$, we have

$$\mathrm{L}(x) \overset{d}{=} \mathcal{S}(\psi(x)) \quad ,$$

where $\psi : \mathcal{X} \to \mathcal{X} \cup \{\perp\}$ — the *oracle* — is such that

$$\psi(x) = \begin{cases} x, \text{ with probability } \epsilon_x \in [0,1], \\ \perp, \text{ otherwise} \end{cases} \quad . \tag{5}$$

We say that L is *ϵ-ARP-simulatable* if there exists a tuple $(\epsilon_x)_{\mathcal{X}}$ with $\mathop{\mathbb{E}}_x[\epsilon_x] = \epsilon$,[7] for which L is simulatable in the $(\epsilon_x)_{\mathcal{X}}$-average random probing. Likewise, we say that L is *ARP-simulatable* if there exists $(\epsilon_x)_{\mathcal{X}}$ with at least two entries strictly lower than 1,[8] for which L is simulatable in the $\{\epsilon_x\}_x$-average random probing model.

Definition 5 is the definition on which Brian *et al.* rely [BDF24]. Contrary to a simulator in the random probing model, where the φ function is fixed, one (*e.g.*, an adversary) may specify the desired probabilities ϵ_x to define the ψ function when building a corresponding simulator. Hereafter, we provide the definition of Dziembowski *et al.* [DFS15b].

Definition 6 (DFS-Average Random Probing)**.** A noisy function $\mathrm{L} : \mathcal{X} \to \mathcal{L}$ is said to be *simulatable in the $(\epsilon_x)_{\mathcal{X}}$-DFS-average random probing model* if there exists a randomized function $\mathcal{S} : \mathcal{X} \cup \{\perp\} \times \mathcal{R} \to \mathcal{L}$ — the *simulator* — such that for every $x \in \mathcal{X}$, we have

$$\mathrm{L}(x) \overset{d}{=} \mathcal{S}(\psi(x, \mathrm{R}), \mathrm{R}) \quad ,$$

[6] Notice that the quantity $\min_x \frac{\Pr(\mathrm{X}=x|\mathrm{L}=l)}{\Pr(\mathrm{X}=x)}$ is upper bounded by one. Indeed, if there is one value x such that $\Pr(\mathrm{X} = x \mid \mathrm{L} = l) > \Pr(\mathrm{X} = x)$ then there is necessarily another value x' such that $\Pr(\mathrm{X} = x' \mid \mathrm{L} = l) > \Pr(\mathrm{X} = x')$, otherwise the PMF $\Pr(\mathrm{X} \mid \mathrm{L} = l)$ cannot sum to one.

[7] Here x is assumed to be uniform over $\mathbb{F}$.

[8] If not, the function ψ is injective (see Definition 7) and the simulation becomes trivial.

where $\psi : \mathcal{X} \times \mathcal{R} \to \mathcal{X} \cup \{\perp\}$ — the *oracle* — verifies Equation 5. Here, $\mathcal{R}$ is any set over which some internal randomness R may be drawn.

The only difference between Definition 5 and Definition 6 is that an adversary in the DFS-ARP model has access to some extra information, namely the internal randomness R used by the oracle to return $\psi(x, \mathrm{R})$. As a result, it is trivial to simulate the view of an adversary in the ARP model from the view of an adversary in the DFS-ARP model, while keeping the $(\epsilon_x)_{\mathcal{X}}$ parameters unchanged. This implies the following bound on the statistical bias of the ARP.

Lemma 2. *The statistical bias of the* ARP *model is upper bounded by* $\delta_{\mathsf{ARP}} = 2\epsilon_{\mathsf{ARP}}$.

Proof. The ARP is trivially simulatable by the DFS-ARP, and since the statistical bias verifies the data processing inequality (Proposition 1), we have $\delta_{\mathsf{ARP}} \leq \delta_{\mathsf{DFS\text{-}ARP}}$. Dziembowski *et al.* have shown that $\delta_{\mathsf{DFS\text{-}ARP}} \leq 2\epsilon_{\mathsf{DFS\text{-}ARP}}$ [DFS15b, Lemma 8]. We conclude by observing that in the simulation, $\epsilon_{\mathsf{DFS\text{-}ARP}} = \epsilon_{\mathsf{ARP}}$. □

4.1 The DFS Simulator

Hereupon, Dziembowski *et al.* have shown that any δ-noisy leakage function is δ-DFS-ARP-simulatable. We transcript hereafter their simulator (denoted DFS thereafter), such as it is described in their paper — up to some change of notation for consistency.

- First, the simulator samples "offline (*i.e.*, independently of the 'real' x) a value l according to the distribution of L(X)" [DFS15b, p. 172].
- Second, $\psi(x, l)$ returns "$\perp$" to the simulator $\mathcal{S}$ with probability $\min\left\{1, \frac{\Pr(\mathrm{L}(x)=l)}{\Pr(\mathrm{L}(\mathrm{X})=l)}\right\}$, and x otherwise.
- Third, the simulator $\mathcal{S}$ returns the following value, depending on its input:
  - Upon receiving "$\perp$", the simulator returns the leakage l sampled at first.
  - Otherwise, upon receiving "x", the simulator re-samples another leakage l' according to the distribution

$$\max\left\{0, \frac{\Pr\left(\mathrm{L}(x) = l'\right) - \Pr\left(\mathrm{L}(\mathrm{X}) = l'\right)}{\mathsf{TV}\left(\mathrm{L}(x); \mathrm{L}(\mathrm{X})\right)}\right\} .$$

One can verify that in the simulator described above, upon receiving x as fixed input, ψ returns x with probability

$$\epsilon_x = \mathop{\mathbb{E}}_{l}\left[1 - \min\left\{1, \frac{\Pr\left(\mathrm{L}(x) = l\right)}{\Pr\left(\mathrm{L}(\mathrm{X}) = l\right)}\right\}\right] = \mathsf{TV}\left(\Pr\left(\mathrm{L}(x)\right); \Pr\left(\mathrm{L}(\mathrm{X})\right)\right),$$

which is confirmed by [DFS15b, Eq. (20)]. Hence, $\epsilon_{\mathsf{DFS}} = \mathop{\mathbb{E}}_{x}\left[\mathsf{TV}\left(\Pr\left(\mathrm{L}(x)\right); \Pr\left(\mathrm{L}(\mathrm{X})\right)\right)\right] = \Delta\left(\mathrm{L}; \mathrm{X}\right)$, which suits well the authors' hope to get rid of the field-size factor.

4.2 The Issue with the Offline Sampling

We now argue why the DFS simulator cannot work in the reduction of noisy leakage to ARP. To build an appropriate simulator in the ARP model, one must verify the same conditions listed in the proof of Theorem 1. The only difference is that nowadays, every ϵ should be replaced by ϵ_x, and is allowed to depend on x.

Unfortunately, despite this simulator seems to verify all the conditions listed in the proof of Theorem 1, it does not verify item 2c, namely that upon receiving "$\perp$", the

leakage distribution of the simulator should not depend on any input. To see why, observe that despite the leakage returned in that case has been initially sampled "offline", it still implicitly relies on the input x, as the second step of the simulator acts as a (soft) rejection sampling, depending on $\Pr(\mathrm{L}(x) = l)$. This is testified in the proof of [DFS15b, Lemma 7]:

> Since "$\perp$" indicates that l [...] is "correct for the real x", in this case [the simulator] simply outputs l [DFS15b, pp.172, 173].

More precisely, on the one hand, the authors prove that for any input $x \in \mathcal{X}$, and for any leakage $l \in \mathcal{L}$, the following quantity — (14) in [DFS15b]:

$$\Pr(\psi(x, \mathrm{R}) = \perp, \mathcal{S}(\psi(x, \mathrm{R})) = l)$$

is equal to $\min\{\Pr(\mathrm{L}(x) = l), \Pr(\mathrm{L}(\mathrm{X}) = l)\}$. On the other hand, they prove that the quantity $\Pr(\psi(x, \mathrm{R}) \neq \perp)$ — (16) in [DFS15b] — is equal to $\mathsf{TV}(\mathrm{L}(x); \mathrm{L}(\mathrm{X}))$ — (20) in [DFS15b]. Using these two equalities, one can notice that[9]

$$\begin{aligned}\Pr(\mathcal{S}(\perp, \mathrm{R}) = l) &= \frac{\min\{\Pr(\mathrm{L}(x) = l), \Pr(\mathrm{L}(\mathrm{X}) = l)\}}{1 - \mathsf{TV}(\mathrm{L}(x); \mathrm{L}(\mathrm{X}))} \\ &= \frac{\min\{\Pr(\mathrm{L}(x) = l), \Pr(\mathrm{L}(\mathrm{X}) = l)\}}{\sum_{l \in \mathcal{L}} \min\{\Pr(\mathrm{L}(x) = l), \Pr(\mathrm{L}(\mathrm{X}) = l)\}} \ ,\end{aligned}$$

which does depend on x. As an example, if we take $\mathrm{L}(x)$ as the LSB of x in a field of size 2^n, one can show that $\mathsf{TV}(\mathrm{L}(x); \mathrm{L}(\mathrm{X})) = \frac{1}{2}$ for any x, but if $l = 0$, we get that

$$\begin{aligned}\min\{\Pr(\mathrm{L}(0) = 0), \Pr(\mathrm{L}(\mathrm{X}) = 0)\} &= \min\left\{1, \frac{1}{2}\right\} = \frac{1}{2} \ , \\ \min\{\Pr(\mathrm{L}(1) = 0), \Pr(\mathrm{L}(\mathrm{X}) = 0)\} &= \min\left\{0, \frac{1}{2}\right\} = 0 \ .\end{aligned}$$

Hence it breaks the condition stated in item 2c.

5 The Average Probing is not Tight with Noisy Leakage

Now that we have emphasized the issue with the DFS simulator, we may naturally wonder whether there exists some correct simulator in the average random probing — beyond the trivial one exhibited in the proof of Theorem 1; and if so, to which extent the ϵ_{ARP} parameter is affected, compared to $\epsilon_{\mathsf{DFS\text{-}ARP}}$. We answer positively to the first question, by showing that any (non-injective) noisy leakage is ARP-simulatable, in a non-trivial manner. However, we answer negatively to the second question, by exhibiting some δ-noisy leakage model that is ϵ_{ARP}-ARP-simulatable, yet with $\epsilon_{\mathsf{ARP}} = \Omega(|\mathcal{X}| \cdot \delta)$.

5.1 A Necessary and Sufficient Condition for ARP-Simulation

To prove these two claims, we first derive a necessary and sufficient condition for simulation in the average random probing.

Lemma 3. *Let $\mathrm{L} : \mathcal{X} \to \mathcal{L}$ be a leakage function. Let $(\epsilon_x)_{\mathcal{X}}$, and define $\mathcal{X}' = \{x \in \mathcal{X} : \epsilon_x < 1\}$. Then, L is simulatable in the $(\epsilon_x)_{\mathcal{X}}$-average random probing model* if and only if *(1) $\mathcal{X}'$ is not empty, and (2) the following inequality holds:*

$$1 \leq \sum_{l \in \mathcal{L}} \min_{x \in \mathcal{X}'} \left\{ \frac{\Pr(\mathrm{L}(x) = l)}{1 - \epsilon_x} \right\} \ . \tag{6}$$

Proof. We proceed by showing both implications.

[9] We use here that $1 - \mathsf{TV}(\mathrm{L}(x); \mathrm{L}(\mathrm{X})) = \sum_{l \in \mathcal{L}} \min\{\Pr(\mathrm{L}(x) = l), \Pr(\mathrm{L}(\mathrm{X}) = l)\}$.

$\Longrightarrow$: Suppose L is simulatable in the $(\epsilon_x)_{\mathcal{X}}$-average random probing model. Condition (1) is trivial, since otherwise $\epsilon_x = 1$ for all input $x \in \mathcal{X}$. In order to prove condition (2), one may simply re-play the proof of Theorem 1, as it should verify the same constraints. The only difference is that now ϵ_x is allowed to depend on x, so the quantity $\pi(l)$ nowadays depends on x and shall be denoted by $\pi(l, x)$ hereafter. As a result, summing Equation 2 becomes

$$\sum_{l \in \mathcal{L}} \pi(l, x) = 1 - \epsilon_x \ ,$$

for all $x \in \mathcal{X}'$, and Equation 3 becomes

$$\Pr\left(\mathcal{S}\left(\perp\right) = l\right) = \frac{\pi(l, x)}{\sum_{l \in \mathcal{L}} \pi(l, x)} = \frac{\pi(l, x)}{1 - \epsilon_x} \ ,$$

for all $x \in \mathcal{X}'$. Likewise, the inequality $\pi(l, x) \leq \Pr\left(\mathrm{L}(x) = l\right)$ — derived from Equation 4 and the fact that $\Pr\left(\mathcal{S}\left(x\right) = l\right) \geq 0$ for all x — remains valid, so for all $x \in \mathcal{X}'$, $\Pr\left(\mathcal{S}\left(\perp\right) = l\right) \leq \frac{\Pr(\mathrm{L}(x)=l)}{1-\epsilon_x}$. Since $\Pr\left(\mathcal{S}\left(\perp\right) = l\right)$ should still remain independent of any x, the latter inequality is equivalent to

$$\Pr\left(\mathcal{S}\left(\perp\right) = l\right) \leq \min_{x' \in \mathcal{X}'} \left\{ \frac{\Pr\left(\mathrm{L}(x') = l\right)}{1 - \epsilon_{x'}} \right\} \ .$$

So we conclude that $\pi(l, x)$ must verify the following constraint:

$$\pi(l, x) \leq (1 - \epsilon_x) \cdot \min_{x' \in \mathcal{X}'} \left\{ \frac{\Pr\left(\mathrm{L}(x') = l\right)}{1 - \epsilon_{x'}} \right\} \ . \tag{7}$$

Summing Equation 7 over $\mathcal{L}$ gives Equation 6.

$\Longleftarrow$: Suppose that $\mathcal{X}'$ is non-empty and Equation 6 is verified. We may build a simulator. To this end, consider the following randomized algorithm $\mathcal{S} : \mathcal{X} \cup \{\perp\} \to \mathcal{L}$:

- Upon receiving "$\perp$" from the oracle $\psi(x)$, $\mathcal{S}$ returns a leakage l sampled according to the distribution

$$\Pr\left(\mathcal{S}\left(\perp\right) = l\right) = \frac{\min_{x' \in \mathcal{X}'} \left\{ \frac{\Pr\left(\mathrm{L}(x') = l\right)}{1 - \epsilon_{x'}} \right\}}{\sum_{l' \in \mathcal{L}} \min_{x' \in \mathcal{X}'} \left\{ \frac{\Pr(\mathrm{L}(x') = l')}{1 - \epsilon_{x'}} \right\}} \ .$$

- Upon receiving "x", from the oracle $\psi(x)$, $\mathcal{S}$ returns samples a leakage according to the following distribution:
  - If $x \in \mathcal{X}'$:

$$\Pr\left(\mathcal{S}\left(x\right) = l\right) = \frac{\Pr\left(\mathrm{L}(x) = l\right) - (1 - \epsilon_x) \cdot \min_{x' \in \mathcal{X}'} \left\{ \frac{\Pr\left(\mathrm{L}(x') = l\right)}{1 - \epsilon_{x'}} \right\}}{1 - \sum_{l' \in \mathcal{L}} (1 - \epsilon_x) \cdot \min_{x' \in \mathcal{X}'} \left\{ \frac{\Pr(\mathrm{L}(x') = l')}{1 - \epsilon_{x'}} \right\}} \ .$$

  - If $x \notin \mathcal{X}'$: $\Pr\left(\mathcal{S}\left(x\right) = l\right) = \Pr\left(\mathrm{L}(x) = l\right)$, for all $l \in \mathcal{L}$.

Let us show that all the conditions of an appropriate simulator are verified. It is clear from its definition that $\Pr\left(\mathcal{S}\left(\perp\right)\right)$ does not depend on any input $x \in \mathcal{X}$, so item 2c holds. Moreover, for all $l \in \mathcal{L}$ and for all $x \in \mathcal{X}'$, $\Pr\left(\mathcal{S}\left(\perp\right) = l\right) = \frac{\pi(l,x)}{1-\epsilon_x}$, and

$\Pr(\mathcal{S}(x)=l) = \frac{\Pr(\mathrm{L}(x)=l)-\pi(l,x)}{\epsilon_x}$, where $\pi(l,x) = (1-\epsilon_x)\cdot\min_{x'\in\mathcal{X}'}\left\{\frac{\Pr(\mathrm{L}(x')=l)}{1-\epsilon_{x'}}\right\}$.[10] This is sufficient to argue that $\Pr(\mathcal{S}(\perp))$ and $\Pr(\mathcal{S}(x))$ are probability mass functions — this trivially holds as well if $x \notin \mathcal{X}'$. Thus, the conditions stated at item 1 and item 2 hold. Finally, using the total probability formula,

$$\begin{aligned}\Pr(\mathcal{S}(\psi(x))=l) =& \epsilon_x\cdot\Pr(\mathcal{S}(x)=l) && +(1-\epsilon_x)\cdot\Pr(\mathcal{S}(\perp)=l) \ ,\\ =& \epsilon_x\cdot\frac{\Pr(\mathrm{L}(x)=l)-\pi(l,x)}{\epsilon_x} && +(1-\epsilon_x)\cdot\frac{\pi(l,x)}{1-\epsilon_x} \ ,\\ =& \Pr(\mathrm{L}(x)=l)-\pi(l,x) && +\pi(l,x)\\ =& \Pr(\mathrm{L}(x)=l) \ .\end{aligned}$$

Hence, item 3 is verified. □

Remark 2. Equation 6 can be seen as a relaxation of the condition of Theorem 1. Indeed, when further assuming ϵ_x to be constant with respect to x, *i.e.*, $\epsilon_x = \epsilon < 1$ for all $x \in \mathcal{X}$, we find back the condition $\epsilon \geq 1 - \sum_l \min_x \Pr(\mathrm{L}(x)=l)$.

We have derived a necessary and sufficient condition on the ϵ_x parameters of the oracle for simulation in the average random probing model. However, contrary to the necessary and sufficient condition for simulation in the random probing, stated in Theorem 1, Equation 6 is not an explicit constraint on the $(\epsilon_x)_{\mathcal{X}}$. The next lemma shows though that for the optimal simulation, Equation 6 becomes an equality.

Lemma 4. *Let* $\mathcal{E} = \{(\epsilon_x)_{\mathcal{X}} : (\epsilon_x)_{\mathcal{X}}$ *verifies Equation 6*$\}$*, and let* $(\epsilon^\star_x)_{\mathcal{X}}$ *minimizing* $\mathop{\mathbb{E}}_x[\epsilon_x]$ *over* $\mathcal{E}$*. Then, it holds that*

$$1 = \sum_{l\in\mathcal{L}} \min_{x\in\mathcal{X}'}\left\{\frac{\Pr(\mathrm{L}(x)=l)}{1-\epsilon^\star_x}\right\} \ .$$

Proof. Let $(\epsilon^\star_x)_{\mathcal{X}}$ be the global minimum of $\mathop{\mathbb{E}}_x[\epsilon_x]$, over $\mathcal{E}$, and define $\mathcal{X}' = \{x \in \mathcal{X} : \epsilon^\star_x < 1\}$. Since $(\epsilon^\star_x)_{\mathcal{X}}$ is a global minimum, it is also a local minimum, *i.e.*, there is a neighbourhood $\mathcal{N} \subset [0,1)^{\mathcal{X}'}$ around the restriction of $(\epsilon^\star_x)_{\mathcal{X}'}$ to $\mathcal{X}'$. This means that $(\epsilon^\star_x)_{\mathcal{X}'}$ should also minimize $\sum_{x\in\mathcal{X}'}\epsilon_x$, subject to $f((\epsilon^\star_x)_{\mathcal{X}'}) \geq 1$, where

$$f : \begin{cases}\mathcal{N} \subset [0,1)^{\mathcal{X}'} \to \mathbb{R}^+\\ (\epsilon_x)_{\mathcal{X}'} \mapsto \sum_l \min_{x\in\mathcal{X}'}\frac{\mathsf{p}_{l,x}}{1-\epsilon_x}\end{cases} \ .$$

Note that f is continuous in $(\epsilon^\star_x)_{\mathcal{X}'}$. Furthermore, f is monotonically increasing in each variable when the others are fixed. We show that necessarily $f((\epsilon^\star_x)_{\mathcal{X}'}) = 1$. By contradiction assume that $f((\epsilon^\star_x)_{\mathcal{X}'}) > 1$. We distinguish two cases.

If $\epsilon^\star_x = 0$ for all $x \in \mathcal{X}'$. Then, $f((\epsilon^\star_x)_{\mathcal{X}'}) = \sum_l \min_{x\in\mathcal{X}'}\mathsf{p}_{l,x} \leq 1$, with equality if and only if the leakage is deterministic (see Subsection 6.2). This contradicts $f((\epsilon^\star_x)_{\mathcal{X}'}) > 1$, so we may exclude this case.

If $\epsilon^\star_x > 0$ for some $x \in \mathcal{X}'$. Let $x \in \mathcal{X}'$ be such that $\epsilon^\star_x > 0$. Then, by continuity and monotony of f, there exists some $0 < a < \epsilon^\star_x$ such that $(\tilde{\epsilon}_x)_{\mathcal{X}'} \in \mathcal{N}$ defined by $\tilde{\epsilon}_x = a$ and $\tilde{\epsilon}_{x'} = \epsilon^\star_{x'}$ for any $x' \in \mathcal{X}'$ not equal to x, is such that $f((\epsilon^\star_x)_{\mathcal{X}'}) \geq f((\tilde{\epsilon}_x)_{\mathcal{X}'}) \geq 1$. But $\sum_{x\in\mathcal{X}'}\tilde{\epsilon}_x < \sum_x \epsilon^\star_x$ which contradicts the minimality of $(\epsilon^\star_x)_{\mathcal{X}'}$.

Hence we can conclude that $f((\epsilon^\star_x)_{\mathcal{X}'}) = 1$. □

[10] If $\epsilon_x = 0$ for some input $x \in \mathcal{X}$, then one does not need to define $\Pr(\mathcal{S}(x))$, since "x" is never passed to the simulator.

A First Limitation of the ARP. We conclude this subsection on the characterization of the ARP model by deriving a first negative consequence of Lemma 3. As already shown in Subsection 4.2, the DFS simulator cannot be used to simulate any δ-noisy leakage in the $(\delta_x)_{\mathcal{X}}$-ARP, where $\delta_x = \mathsf{TV}\left(\mathrm{L}(x); \mathrm{L}(\mathrm{X})\right)$ — whereas it was possible in the DFS-ARP model as explained in Subsection 4.2. To this end, we used the lsb leakage model as a counter-example. With Equation 6 nowadays, we may revisit this use case to prove an even stronger negative result: there is *no* simulator for which the $(\delta_x)_{\mathcal{X}}$-ARP-simulatability holds.[11]

Corollary 1. *Let* lsb *be the function returning the least significant bit of its input. Then, for any* $x \in \mathcal{X}$, $\mathsf{TV}\left(\mathsf{lsb}\left(x\right); \mathsf{lsb}\left(\mathrm{X}\right)\right) = \frac{1}{2}$, *but* lsb *is not simulatable in the* $\left(\frac{1}{2}, \frac{1}{2}, \ldots, \frac{1}{2}\right)$*-average random probing model.*

Proof. Observe that we have $\delta_x = \mathsf{TV}\left(\mathsf{lsb}\left(x\right); \mathsf{lsb}\left(\mathrm{X}\right)\right) = \frac{1}{2}$ for any $x \in \mathcal{X}$. It implies that $\mathcal{X}' = \mathcal{X}$, yet for any leakage $l \in \{0, 1\}$, there exits some input x such that $\Pr\left(\mathsf{lsb}\left(x\right) = l\right) = 0$. Hence, Equation 6 is not verified. □

5.2 Universality of the Average Random Probing

Provided with the new necessary and sufficient condition of Lemma 3, we may now address the first question raised at the beginning of Section 5, namely the fact that any *non-injective* leakage function is ARP-simulatable. By "non-injective", we mean a leakage function not verifying the following definition.

Definition 7 (Injectivity)**.** A (possibly randomized) leakage function $\mathrm{L} : \mathcal{X} \to \mathcal{L}$ is said to be *injective* if for any leakage $l \in \mathcal{L}$, there is a unique x such that $\Pr\left(\mathrm{L}(x) = l\right) > 0$.

Theorem 2 (ARP-Simulatability)**.** *Any non-injective leakage function is ARP-simulatable. Reciprocally, any ARP-simulatable leakage is non-injective.*

Proof. Let $\mathrm{L} : \mathcal{X} \to \mathcal{L}$ be a leakage function. We proceed by both implication.

$\Longleftarrow$: Suppose that L is injective. We first show that the set $\mathcal{X}'$ of values for which $\epsilon_x < 1$ is reduced to a singleton. By contradiction, assume that there exists at least two values $x_0, x_1 \in \mathcal{X}$ such that $\epsilon_{x_0} < 1$ and $\epsilon_{x_1} < 1$. We may then upper bound the right-hand side of Equation 6 as follows:

$$\sum_{l \in \mathcal{L}} \min_{x \in \mathcal{X}'} \left\{ \frac{\Pr\left(\mathrm{L}(x) = l\right)}{1 - \epsilon_x} \right\} \leq \sum_{l \in \mathcal{L}} \min \left\{ \frac{\Pr\left(\mathrm{L}(x_0) = l\right)}{1 - \epsilon_{x_0}}, \frac{\Pr\left(\mathrm{L}(x_1) = l\right)}{1 - \epsilon_{x_1}} \right\}.$$

For all term in the latter sum, according to Definition 7, if $\Pr\left(\mathrm{L}(x_0) = l\right) > 0$, then necessarily $\Pr\left(\mathrm{L}(x_1) = l\right) = 0$, and inversely. In other words, every term of the sum is upper bounded by 0, which transgresses the condition of Equation 6. Therefore, the set $\mathcal{X}'$ is reduced to one singleton $\{x_0\}$ at most, with $\epsilon_{x_0} = 1 - \Pr\left(\mathrm{L}(x_0) \in \mathcal{L}_{x_0}\right) = 0$, in order to verify Equation 6. But this would result in a trivial simulation where the oracle always returns its input, by only substituting "x_0" with "$\perp$" in the alphabet $\mathcal{X}$.

$\Longrightarrow$: Suppose that L is non-injective. Then there exists a leakage $l^\star$ and at least two inputs $x_0, x_1 \in \mathcal{X}$ such that $0 < \Pr\left(\mathrm{L}(x_0) = l^\star\right)$ and $0 < \Pr\left(\mathrm{L}(x_1) = l^\star\right)$. Define $\epsilon_{x_0} = 1 - \Pr\left(\mathrm{L}(x_0) = l^\star\right), \epsilon_{x_1} = 1 - \Pr\left(\mathrm{L}(x_1) = l^\star\right)$, and $\epsilon_x = 1$ for any $x \notin \{x_0, x_1\}$. The sum in Equation 6 may be simplified as

$$\sum_{l \in \mathcal{L}} \min_{x \in \mathcal{X}'} \frac{\Pr\left(\mathrm{L}(x) = l\right)}{1 - \epsilon_x} \geq \min \left\{ \frac{\Pr\left(\mathrm{L}(x_0) = l^\star\right)}{1 - \epsilon_{x_0}}, \frac{\Pr\left(\mathrm{L}(x_1) = l^\star\right)}{1 - \epsilon_{x_1}} \right\} = 1 \ .$$

[11] We will actually nuance this negative result in Subsection 6.2 showing that the lsb leakage model remains $\frac{1}{2}$-ARP-simulatable, despite not being $\left(\frac{1}{2}, \frac{1}{2}, \ldots, \frac{1}{2}\right)$-ARP-simulatable.

Therefore, the necessary and sufficient condition of Equation 6 is verified. □

A similar result holds for the DFS-ARP model, since it is essentially equivalent to the noisy leakage model [DFS15b, Lemma 2]. On the opposite, the ARP contrasts with the RP in that respect, since Theorem 2 would not extend for the latter leakage model.

5.3 A Catastrophic Channel for Average Random Probing

We now address the second question raised in Section 5, namely to which extent the reduction is tight. To this end, consider the following side channel

$$\mathrm{X} \to \boxed{\mathsf{p}_{\mathrm{L}\mid\mathrm{X}}} \to \mathrm{L}$$

where $\mathcal{X} = \mathcal{L}$ and $\Pr(\mathrm{L}(x) = l) \triangleq \frac{1_{x\neq l}}{|\mathcal{X}|-1}$. In other words, observing $\mathrm{L} = l$ means that the adversary receives the information "x cannot take the value l".[12]

It can be verified, on the one hand, that for the catastrophic channel, $\Delta(\mathrm{L};\mathrm{X}) = |\mathcal{X}|^{-1}$. On the other hand, the following proposition states that the optimal parameter of the ARP-simulation verifies $\epsilon_{\mathsf{ARP}} = \Omega(1) = \Omega(\delta \cdot |\mathcal{X}|)$.

Proposition 3. *Let* $\mathrm{L} : x \mapsto l$ *with probability* $\Pr(\mathrm{L}(x) = l) = \frac{\mathbf{1}_{x\neq l}}{|\mathcal{X}|-1}$. *Let* $\epsilon \in [0,1]$ *be any parameter such that* L *is* ϵ*-ARP-simulatable. Then,*

$$\epsilon \geq 1 - \frac{1}{4} \cdot \frac{|\mathcal{X}|}{|\mathcal{X}|-1} \geq \frac{1}{2}.$$

Proof. Let $(\epsilon_x)_{\mathcal{X}}$ be any set of parameters such that L is $(\epsilon_x)_{\mathcal{X}}$-ARP-simulatable, and define $\mathcal{X}' = \{x \in \mathcal{X} : \epsilon_x < 1\}$. For any $l \in \mathcal{X}$, if $l \in \mathcal{X}'$ then

$$\min_{x\in\mathcal{X}'} \frac{\Pr(\mathrm{L}(x) = l)}{1-\epsilon_x} = \frac{\Pr(\mathrm{L}(l) = l)}{1-\epsilon_l} = 0.$$

It results that the necessary condition for ARP-simulatability stated in Equation 6 may be rephrased as

$$1 \leq \sum_{l\notin\mathcal{X}'} \min_{x\in\mathcal{X}'} \frac{\Pr(\mathrm{L}(x) = l)}{1-\epsilon_x} = \sum_{l\notin\mathcal{X}'} \min_{x\in\mathcal{X}'} \frac{\frac{1}{|\mathcal{X}|-1}}{1-\epsilon_x} = \frac{|\mathcal{X}| - |\mathcal{X}'|}{|\mathcal{X}|-1} \cdot \frac{1}{1-\min_{x\in\mathcal{X}'}\epsilon_x},$$

where the first inequality comes from Equation 6, the first equality holds by definition of $\Pr(\mathrm{L}(x) = l)$ is our case, and the second equality follows from the fact that for any $x \neq l$, the quantity $\Pr(\mathrm{L}(x) = l)$ is constant. As a result, we have

$$\min_{x\in\mathcal{X}'} \epsilon_x \geq 1 - \frac{|\mathcal{X}| - |\mathcal{X}'|}{|\mathcal{X}|-1}. \tag{8}$$

It now remains to compute the average over all ϵ_x:

$$\begin{aligned}
\epsilon = \mathop{\mathbb{E}}_{x}[\epsilon_x] &= \frac{|\mathcal{X}'|}{|\mathcal{X}|} \cdot \mathop{\mathbb{E}}_{x\in\mathcal{X}'}[\epsilon_x] + \frac{|\mathcal{X}| - |\mathcal{X}'|}{|\mathcal{X}|} \cdot \mathop{\mathbb{E}}_{x\notin\mathcal{X}'}[\epsilon_x] \\
&\geq \frac{|\mathcal{X}'|}{|\mathcal{X}|} \cdot \min_{x\in\mathcal{X}'} \epsilon_x + \frac{|\mathcal{X}| - |\mathcal{X}'|}{|\mathcal{X}|} \cdot 1 \\
&\geq \frac{|\mathcal{X}'|}{|\mathcal{X}|} \cdot \left(1 - \frac{|\mathcal{X}| - |\mathcal{X}'|}{|\mathcal{X}|-1}\right) + \frac{|\mathcal{X}| - |\mathcal{X}'|}{|\mathcal{X}|} \\
&= 1 - \frac{|\mathcal{X}'|}{|\mathcal{X}| \cdot (|\mathcal{X}|-1)} \cdot (|\mathcal{X}| - |\mathcal{X}'|),
\end{aligned}$$

[12] Interestingly, this leakage function has already been considered by Brian *et al.* to show some limitations when reducing the noisy leakage model to the so-called *bounded leakage* model [BFO+21, Thm.10, Eq. (13)].

where the first inequality is obtained by lower-bounding the expectation over $\mathcal{X}'$ by the minimum over $\mathcal{X}'$, and the second inequality holds by virtue of Equation 8. The last right-hand side, seen as a function of $|\mathcal{X}'|$, takes its minimum value for $|\mathcal{X}'| = \frac{1}{2} \cdot |\mathcal{X}|$, which gives $\epsilon \geq 1 - \frac{1}{4} \cdot \frac{|\mathcal{X}|}{|\mathcal{X}|-1} \geq \frac{1}{2}$. □

In other words, this example proves that there are some δ-noisy leakage functions $\mathrm{L} : \mathcal{X} \to \mathcal{L}$ that are ϵ-ARP-simulatable, although at the condition that $\epsilon = \Omega(\delta \cdot |\mathcal{X}|)$. This contradicts the claim that noisy leakage can be simulated from the ARP without the field-size loss, as it is feasible from the DFS-ARP.

6 When is the Reduction Tight?

As seen in Section 5, it is not generally possible to have a tight reduction, *i.e.*, without field-size loss, from the noisy leakage model to the average random probing model, since we have exhibited a counter-example. Nevertheless, that does not mean that the reduction is loose for *any* leakage function. In this subsection, we emphasize two classes of leakage functions for which a finer characterization, beyond measuring the noisy leakage parameter δ, may allow for a tight reduction.

The former one concerns leakage functions whose marginal distribution has a bounded, *compact* support. It covers very particular cases but for which the noisy leakage parameter δ may be arbitrarily low. The latter one deals with deterministic leakage functions. In that respect, there is no restriction on the leakage range, but by definition, it only covers a few use cases such that $\delta \geq |\mathcal{X}|^{-1}$.[13]

6.1 Leakage Functions with Compact Support

In this first use-case, we consider leakage functions such that the marginal distribution is *compact*, in the sense that the least likely observable leakage $l^\star$ is such that $\Pr(\mathrm{L} = l^\star) \gg |\mathcal{X}|^{-1}$. This class of leakages is contained into the class of bounded leakage functions, although not all bounded leakage functions are compact. This is formalized by the following proposition.

Proposition 4. *Let* $\mathrm{L} : \mathcal{X} \to \mathcal{L}$ *be a* δ*-noisy leakage function. Then,* L *is* ϵ*-ARP-simulatable, with*

$$\epsilon = \frac{\delta}{\min_{l \in \mathcal{L}} \Pr(\mathrm{L} = l)} \ .$$

Proof. Let T be any random variable with full support over $\mathcal{L}$. For all $x \in \mathcal{X}$, define

$$\epsilon_x = 1 - \min_{l \in \mathcal{L}} \frac{\Pr(\mathrm{L}(x) = l)}{\Pr(\mathrm{T} = l)} \ .$$

Then for any $l \in \mathcal{L}$, we have that $\frac{1}{1-\epsilon_x} \geq \frac{\Pr(\mathrm{T}=l)}{\Pr(\mathrm{L}(x)=l)}$ so for all $x \in \mathcal{X}'$:

$$\frac{\Pr(\mathrm{L}(x) = l)}{1 - \epsilon_x} \geq \Pr(\mathrm{T} = l) \ .$$

Hence,

$$\sum_{l \in \mathcal{L}} \min_{x \in \mathcal{X}'} \frac{\Pr(\mathrm{L}(x) = l)}{1 - \epsilon_x} \geq \sum_{l \in \mathcal{L}} \Pr(\mathrm{T} = l) = 1 \ ,$$

[13] To see why, notice that any deterministic leakage function is not RP-simulatable, except the constant function, *i.e.* $\epsilon_{\mathsf{RP}} = 1$. The claim follows from Lemma 1.

i.e., the necessary and sufficient condition of Equation 6 for ARP-simulatability is verified. Let us now compute and bound $\epsilon = \mathop{\mathbb{E}}_{x}[\epsilon_x]$, in the case where $\mathrm{T} = \mathrm{L}(\mathrm{X})$.

$$\epsilon = 1 - \mathop{\mathbb{E}}_{x}\left[\min_{l\in\mathcal{L}} \frac{\Pr(\mathrm{L}(x)=l)}{\Pr(\mathrm{L}(\mathrm{X})=l)}\right] = 1 - \mathop{\mathbb{E}}_{x}\left[\min_{l\in\mathcal{L}} \frac{\Pr(\mathrm{X}=x \mid \mathrm{L}=l)}{\Pr(\mathrm{X}=x)}\right] \tag{9}$$

$$= 1 - \sum_{x\in\mathcal{X}} \min_{l\in\mathcal{L}} \Pr(\mathrm{X}=x \mid \mathrm{L}=l)$$

$$= \sum_{x\in\mathcal{X}} \Pr(\mathrm{X}=x) - \min_{l\in\mathcal{L}} \Pr(\mathrm{X}=x \mid \mathrm{L}=l) \tag{10}$$

$$= \sum_{x\in\mathcal{X}} \max_{l\in\mathcal{L}} \{\Pr(\mathrm{X}=x) - \Pr(\mathrm{X}=x \mid \mathrm{L}=l)\}$$

$$\leq \sum_{x\in\mathcal{X}} \sum_{l\in\mathcal{L}} \max\{0, \Pr(\mathrm{X}=x) - \Pr(\mathrm{X}=x \mid \mathrm{L}=l)\} \tag{11}$$

$$= \sum_{l\in\mathcal{L}} \mathsf{TV}(\Pr(\mathrm{X}); \Pr(\mathrm{X} \mid \mathrm{L}=l)) \tag{12}$$

$$\leq \frac{1}{\min_{l\in\mathcal{L}} \Pr(\mathrm{L}=l)} \cdot \sum_{l\in\mathcal{L}} \Pr(\mathrm{L}=l) \cdot \mathsf{TV}(\Pr(\mathrm{X}); \Pr(\mathrm{X} \mid \mathrm{L}=l)) \tag{13}$$

$$= \frac{\Delta(\mathrm{X};\mathrm{L})}{\min_{l\in\mathcal{L}} \Pr(\mathrm{L}=l)} = \frac{\delta}{\min_{l\in\mathcal{L}} \Pr(\mathrm{L}=l)} \quad .$$

In the latter development, we have used Bayes' theorem to get Equation 9. We have also rephrased 1 as the sum over $\mathcal{X}$ of $\Pr(\mathrm{X}=x)$ in Equation 10. Equation 11 comes from the fact that the maximum of positive values cannot be larger than their sum.[14] Equation 12 holds by definition of the total variation distance – see Section 2. Finally, we get Equation 13 by observing that for any $l \in \mathcal{L}$, we have that $1 \leq \frac{\Pr(\mathrm{L}=l)}{\min_{l'\in\mathcal{L}} \Pr(\mathrm{L}=l')}$. □

This proposition tells us that we may bound the average random probing parameter ϵ by a function of δ and of $\min_{l\in\mathcal{L}} \Pr(\mathrm{L}=l)$. It is worth emphasizing that the choice of T in the proof, and thereby ϵ_x, is not necessarily optimal, *i.e.*, there is a chance that the bound given in Proposition 4 may be improved, *e.g.*, by using the bounded retrieval model with a short description length.

6.2 The Case of Deterministic Leakages

Another interesting use case — perhaps less unrealistic — is the class of deterministic leakage functions. It is easy to see that the leakage functions in this class are not RP-simulatable, since there is no leakage that can be generated from any input, unless for the constant leakage function. Nevertheless, such leakages remain ARP-simulatable by virtue of Theorem 2.

A second reason to the interest in deterministic leakages is that it allows to turn the implicit constraint of Equation 6 into explicit constraints on each ϵ_x, which in turn allows to derive the optimal simulator in those cases. To this end, we may notice that for deterministic leakage functions, any input can only be associated to one leakage value. In other words, the input space can be partitioned into several sets of pre-images.

Lemma 5 (Characterization of Deterministic Leakages)**.** *A leakage function is deterministic if and only if the sets*

$$\mathcal{X}_l = \{x \in \mathcal{X} : \Pr(\mathrm{L}(x)=l) > 0\} \quad ,$$

are mutually exclusive.

[14] See [DDF14, Footnote 8].

Proof. L is deterministic if and only if for any input x, there is only one leakage l such that $\Pr(\mathrm{L}(x) = l) > 0$, so x can only belong to one set $\mathcal{X}_l$. Reciprocally, if the sets are mutually exclusive, any input x belongs to one set $\mathcal{X}_{l^\star}$, so $\Pr(\mathrm{L}(x) = l^\star) = \sum_l \Pr(\mathrm{L}(x) = l) = 1$. □

Theorem 3. *Let* $\mathrm{L} : \mathcal{X} \to \mathcal{L}$ *be a deterministic leakage function. Then* L *is simulatable in the* $(\epsilon_x)_{\mathcal{X}}$*-average random probing model if* and only if *there is some leakage* $l^\star \in \mathcal{L}$ *such that for all* $x \notin \mathcal{X}_{l^\star}$ *we have* $\epsilon_x = 1$.

Proof. The necessary and sufficient condition can be rephrased as

$$\mathcal{X}' = \{x \in \mathcal{X} : \epsilon_x < 1\} \subseteq \mathcal{X}_{l^\star}, \text{ for some } l^\star \in \mathcal{L} \ .$$

We proceed the proof by showing both implications.

$\Longrightarrow$: Suppose $\mathcal{X}' \not\subseteq \mathcal{X}_l$ for all $l \in \mathcal{L}$. Then, for all $l \in \mathcal{L}$, there is some $x \in \mathcal{X}'$ such that $\Pr(\mathrm{L}(x) = l) = 0$. This in turn implies that

$$\sum_{l \in \mathcal{L}} \min_{x \in \mathcal{X}'} \left\{ \frac{\Pr(\mathrm{L}(x) = l)}{1 - \epsilon_x} \right\} = 0 \ ,$$

which breaks the necessary condition of Equation 6.

$\Longleftarrow$: Suppose $\mathcal{X}' \subseteq \mathcal{X}_{l^\star}$ for some $l^\star \in \mathcal{L}$. Since, by assumption, $\{\mathcal{X}_l\}_{l \in \mathcal{L}}$ are mutually exclusive, if $\mathcal{X} \subseteq \mathcal{X}_{l^\star}$ for some leakage $l^\star \in \mathcal{L}$, then such $l^\star$ is unique. Hence,

$$\sum_{l \in \mathcal{L}} \min_{x \in \mathcal{X}'} \left\{ \frac{\Pr(\mathrm{L}(x) = l)}{1 - \epsilon_x} \right\} = \min_{x \in \mathcal{X}'} \left\{ \frac{1}{1 - \epsilon_x} \right\} \geq 1,$$

which verifies Equation 6. In the latter equality, we used the fact that for deterministic leakages, $\Pr(\mathrm{L}(x) = l) > 0 \iff \Pr(\mathrm{L}(x) = l) = 1$.

□

Figure 2 illustrates the simulator, characterized through the proof of Theorem 3: one chooses a single leakage, and only the values belonging to the pre-image set of $l^\star$ are allowed to be revealed to the adversary with some probability ϵ_x strictly less than 1. Given

(a) The leakage function $x \mapsto \mathrm{L}(x)$. (b) The corresponding optimal ARP simulator.

Figure 2: An exemplary deterministic leakage function, and the corresponding optimal simulator.

that Theorem 3 requires no further constraint, for a given fixed $l^\star$, the best simulator — *i.e.*, the one minimizing ϵ — is such that $\epsilon_x = 0$ for any $x \in \mathcal{X}_{l^\star}$. Hence,

$$\epsilon = \mathop{\mathbb{E}}_x [\epsilon_x] = 1 - \frac{|\mathcal{X}_{l^\star}|}{|\mathcal{X}|} \ .$$

Consequently, the optimal simulation is obtained by choosing $l^\star$ as the leakage occurring the most frequently. This allows to conclude with the following corollary.

Corollary 2 (Deterministic Leakages)**.** *Let* $\mathrm{L} : \mathcal{X} \to \mathcal{L}$ *be a deterministic,* δ*-noisy, leakage function. Then,* L *is* δ*-ARP-simulatable.*

Proof. Observe that for deterministic leakages, $\delta = 1 - \frac{\mathbb{E}[|\mathcal{X}_l|]}{|\mathcal{X}|} \geq 1 - \frac{\max_l |\mathcal{X}_l|}{|\mathcal{X}|} = \epsilon$. □

In other words, assuming that each wire leaks a deterministic, δ-noisy function of its value, Corollary 2 confirms that the provably secure construction of Brian *et al.* remains valid [BDF24, Cor. 9].

7 Discussion and Perspectives

We have now exposed the main results of this work, but before concluding this paper, let us synthesize some of its takeaway messages. We have exhibited a gap between the formal definitions given to the average random probing model between Dziembowski *et al.*'s works [DFS15b] and Brian *et al.*'s works [BDF24]. With this clarification in mind, we have shown that contrary to Dziembowski *et al.*'s tight simulation in the DFS-ARP model, the reduction to the ARP defined by Brian *et al.* still conveys a security loss of $\Omega(|\mathcal{X}|)$, questioning the tight reduction from noisy leakage to random probing model claimed by Brian *et al.* Thankfully, this tight reduction remains valid, *e.g.*, if one makes the additional assumption that the leakage function of each wire is deterministic. We begin our discussions hereafter by first elaborating on that point.

7.1 On the Deterministic Leakage Functions

Although far from generic, the deterministic leakage assumption covers physically realistic situations, in a low physical noise setting, *e.g.* when the Hamming weight or whenever any physical bit of each wire is revealed to the adversary. It can also be generalized to the same leakage functions with some independent additive noise, since they can be trivially simulated from their deterministic counterpart. But, the corresponding security bound would be kept unchanged, regardless of this additive noise. This is clearly a proof artifact, since this contradicts the intuition of *noise amplification* of masking [CJRR99].

However, a tight reduction for deterministic leakage functions does not make sound security bounds yet. Framing the discussion in terms of admissible values of δ by the whole chain of security proof, the deterministic-leakage assumption only covers a restricted range. As argued in footnote 13, by definition of deterministic leakage functions, their corresponding noisy-leakage parameter δ must be higher than $|\mathcal{X}|^{-1}$. This constraint must be put into perspectives with two additional constraints: Firstly, Brian *et al.*'s security proof requires $\epsilon_{\mathsf{ARP}} = 35{\epsilon_{\mathsf{RP}}}^2$ [BDF24, Thm. 2]. Secondly, the current RP-secure circuits with a constant tolerated leakage probability require $\epsilon_{\mathsf{RP}} \leq 2^{-7.5}$ [BRT21]. Overall, our patch only works for deterministic leakages verifying

$$|\mathcal{X}|^{-1} \leq \delta \leq \frac{2^{-15}}{35},$$

which in turn implies that $|\mathcal{X}| \geq 2^{20}$. In other words, outside this range, our patch does not improve upon the classical reduction of Duc *et al.* [DDF14]. Such an improvement, conjectured for a decade so far [DFS15a], remains therefore an open problem. Even if the tolerated leakage probability ϵ_{RP} could be increased in a near future, Brian *et al.*'s reduction would still require the optimal ARP parameter for deterministic leakage functions,

namely $\epsilon_{\mathsf{ARP}} = 1 - \max_l \frac{|\mathcal{X}_l|}{|\mathcal{X}|}$, to be at most $\frac{1}{35}$. This *de facto* eliminates most of the physically realistic leakage functions.[15]

7.2 Cautionary Note: How not to Interpret our Result

Although we have shown that in its current state, the construction of Brian *et al.* [BDF24] is not practical, we would like to stress that this does not necessarily invalidate Brian *et al.*'s final result — it only remains unproven yet. In this respect, one must not interpret from our work that the core idea implemented in Brian *et al.*'s construction is a dead-end. Actually, we believe that it remains a promising line of research.

To see why, we need to take a few steps back: so far in this paper, we have considered simulating leakage functions occurring at the scale of a single wire. Proving the reduction at the scale of the whole circuit straightforwardly follows then, as shown by Duc *et al.* [DDF14, Lemma 5]. More precisely, we can show that simulating from the RP at the scale of each wire independently is necessary and sufficient to simulate from the RP at the scale of the whole circuit, for *any* distribution of the circuit wires, as stated hereafter.[16]

Corollary 3 (of [DDF14, Lemma 5])**.** *Let* $\mathrm{L}_1, \ldots, \mathrm{L}_\ell$ *be* ℓ *leakage functions:* $\mathcal{X} \to \mathcal{L}$*, mutually independent, given their respective inputs. Define for short the random vector* $\vec{\mathrm{X}} = (\mathrm{X}_1, \ldots, \mathrm{X}_\ell)$ *for random variables* $\mathrm{X}_i \in \mathcal{X}$*, and* $\mathbf{L}(\vec{\mathrm{X}}) = (\mathrm{L}_1(\mathrm{X}_1), \ldots, \mathrm{L}_\ell(\mathrm{X}_\ell))$*. Then, there exists a simulator* $\boldsymbol{\mathcal{S}} : (\mathcal{X} \cup \{\perp\})^\ell \to \mathcal{L}^\ell$ *and an oracle* $\boldsymbol{\varphi}(\vec{x}) = (\varphi_1(x_1), \ldots, \varphi_\ell(x_\ell))$*, where each* φ_i *is an* ϵ_i*-identity function, such that for* any *distribution over the random variables* $\mathrm{X}_1, \ldots, \mathrm{X}_\ell$*, the following equality holds:*

$$\mathbf{L}(\vec{\mathrm{X}}) \overset{d}{=} \boldsymbol{\mathcal{S}}(\boldsymbol{\varphi}(\vec{\mathrm{X}})) \ , \tag{14}$$

if and only if *each leakage function* L_i *is* ϵ_i*-*RP*-simulatable.*

In a sense, Corollary 3 is what makes the security proof of Duc *et al.* very generic and elegant, as one does not need to bother about the actual distribution of the values in the wires of the circuit. On the other hand, since Corollary 3 is equivalent to Theorem 1, there is no hope to improve the parameters ϵ_i of the RP oracles given by Lemma 1, without any extra assumption.

However, one might not need the reduction from noisy leakage to random probing to hold for *any* distribution over the wires. Instead, a careful scrutiny of the joint distribution of the wires on the masked circuit may allow to restrict the scope of Corollary 3, to some distributions only, for which the parameters $(\epsilon_i)_{1 \le i \le \ell}$ of the oracle might be significantly improved. In a sense, we would like now that for a restricted set of joint distributions p over $\mathcal{X}^\ell$, ϵ may be allowed to depend on p, and not only on some features of the leaky channel L. This is what Brian *et al.* capture through the construction of their circuit compiler, and the subsequent analysis, *e.g.*, by noticing that the joint distribution of some wires is *close to uniform* [BDF24]. In that respect, we still believe that investigating this line of research represents a promising strategy.

7.3 Possible Patches to [BDF24]

We conclude this paper by proposing two ways to circumvent the limitation of the ARP leakage model. A first and natural strategy could be to revisit the analysis conducted by Brian *et al.* based on the ARP model, directly in the DFS-ARP model. Although we did not investigate this line of research, we do not believe this strategy to be promising.

[15] For leakage functions like lsb or Hamming weight, we have $\epsilon_{\mathsf{ARP}} = 1 - \max_l \frac{|\mathcal{X}_l|}{|\mathcal{X}|} \ge \frac{1}{2}$. As an example, attacks are already exhibited from these leakage models if $\mathcal{X}$ is a binary field [MMMS23].

[16] The proof is given in Appendix A.

Indeed, leveraging their reduction from the noisy leakage to the DFS-ARP leakage model, Dziembowski *et al.* have derived security bounds in two use cases: when a single encoding is leaking, and when the Rivain & Prouff compiler [RP10] is applied with leak-free refreshings [PR13]. Although their works cover adaptively-chosen leakage functions – a much stronger, but somewhat unrealistic threat scenario –, the proof techniques used by Dziembowski *et al.* in the DFS-ARP model are arguably tedious. Since then, their bounds have been improved by Béguinot *et al.* [BCG+23, MRS22] and Ito *et al.* [IUH22] in the single-encoding scenario, and by Masure and Standaert [MS23] in the Prouff-Rivain scenario, through direct analysis in the noisy leakage model, with much simpler techniques. This suggests that any security analysis in the DFS-ARP model could be conducted in a simpler way in the noisy leakage model, without requiring Dziembowski *et al.*'s reduction [DFS15b].

Inversely, we might figure out a patch by revisiting Brian *et al.*'s security analysis directly from the noisy leakage model — instead of the ARP. This brings new challenges, as we might lose the *all-or-nothing* nature of the (average) random probing model, that is leveraged in most of the literature on those models. We let this opportunity as an open question for future research.

A Missing Proofs

We first recall Duc *et al.*'s reduction to the random probing model, at the scale of a whole circuit. We then prove that Lemma 6 can be equivalently restated as Corollary 3.

Lemma 6 ([DDF14, Lemma 5], rephrased)**.** *Let* $\mathrm{L}_1, \dots, \mathrm{L}_\ell$ *be* ℓ *leakage functions, mutually independent, given their respective inputs.*[17] *For any* $1 \leq i \leq \ell$*, define* $\epsilon_i = 1 - \sum_{l \in \mathcal{L}} \min_{x \in \mathcal{X}} \Pr\left(\mathrm{L}_i(x) = l\right)$*. Then, there exist* ℓ *oracles* φ_i *— respectively* ϵ_i*-identity functions — and* ℓ *simulators* $\mathcal{S}_1, \dots, \mathcal{S}_\ell$ *such that for any tuple* $x_1, \dots, x_\ell$*,*

$$\mathrm{L}_1(x_1), \dots, \mathrm{L}_\ell(x_\ell) \overset{d}{=} \mathcal{S}_1\left(\varphi_1(x_1)\right), \dots, \mathcal{S}_\ell\left(\varphi_\ell(x_\ell)\right) \quad . \tag{15}$$

Proof of Corollary 3. We prove the equivalence by double implication.

$\Longrightarrow$: Suppose that Equation 15 holds for any tuple $(\boldsymbol{x_1, \dots, x_\ell})$. Using the total probability formula, we have

$$\begin{aligned}
\Pr\left(\mathbf{L}(\vec{\mathrm{X}}) = \mathbf{l}\right) &= \sum_{\vec{x} \in \mathcal{X}^\ell} \Pr\left(\vec{\mathrm{X}} = \vec{x}\right) \cdot \Pr\left(\mathbf{L}(\vec{x}) = \mathbf{l}\right) \\
&= \sum_{\vec{x} \in \mathcal{X}^\ell} \Pr\left(\vec{\mathrm{X}} = \vec{x}\right) \cdot \prod_{i=1}^{\ell} \Pr\left(\mathrm{L}_i(x_i) = l_i\right) && \text{(By indep. given the } x_i\text{s)} \\
&= \sum_{\vec{x} \in \mathcal{X}^\ell} \Pr\left(\vec{\mathrm{X}} = \vec{x}\right) \cdot \prod_{i=1}^{\ell} \Pr\left(\mathcal{S}_i(\varphi_i(x_i)) = l_i\right) && \text{(By Theorem 1)} \\
&= \sum_{\vec{x} \in \mathcal{X}^\ell} \Pr\left(\vec{\mathrm{X}} = \vec{x}\right) \cdot \Pr\left(\boldsymbol{\mathcal{S}}(\boldsymbol{\varphi}(\vec{x})) = \mathbf{l}\right) \\
&= \Pr\left(\boldsymbol{\mathcal{S}}(\boldsymbol{\varphi}(\vec{\mathrm{X}})) = \mathbf{l}\right) && \text{(By total probability)}
\end{aligned}$$

$\Longleftarrow$: Suppose that Equation 14 holds for any distribution of $\mathbf{X_1, \dots, X_\ell}$. In particular, for any tuple $(x_1, \dots, x_\ell)$, this holds whenever $(\mathrm{X}_1, \dots, \mathrm{X}_\ell) = (x_1, \dots, x_\ell)$ with probability 1, hence proving Equation 15. □

[17] *i.e.* for each $i \neq j$ and for any couple x_i, x_j, $\mathrm{L}_i(x_i)$ and $\mathrm{L}_j(x_j)$ are independent.

References

[BBD+16] Gilles Barthe, Sonia Belaïd, François Dupressoir, Pierre-Alain Fouque, Benjamin Grégoire, Pierre-Yves Strub, and Rébecca Zucchini. Strong non-interference and type-directed higher-order masking. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 2016: 23rd Conference on Computer and Communications Security*, pages 116–129, Vienna, Austria, October 24–28, 2016. ACM Press.

[BCG+23] Julien Béguinot, Wei Cheng, Sylvain Guilley, Yi Liu, Loïc Masure, Olivier Rioul, and François-Xavier Standaert. Removing the field size loss from duc et al.'s conjectured bound for masked encodings. In Elif Bilge Kavun and Michael Pehl, editors, *COSADE 2023: 14th International Workshop on Constructive Side-Channel Analysis and Secure Design*, volume 13979 of *Lecture Notes in Computer Science*, pages 86–104, Munich, Germany, April 3–4, 2023. Springer, Cham, Switzerland.

[BCGR24] Julien Béguinot, Wei Cheng, Sylvain Guilley, and Olivier Rioul. Formal security proofs via doeblin coefficients: Optimal side-channel factorization from noisy leakage to random probing. Cryptology ePrint Archive, Paper 2024/199, 2024.

[BCP+20] Sonia Belaïd, Jean-Sébastien Coron, Emmanuel Prouff, Matthieu Rivain, and Abdul Rahman Taleb. Random probing security: Verification, composition, expansion and new constructions. In Daniele Micciancio and Thomas Ristenpart, editors, *Advances in Cryptology – CRYPTO 2020, Part I*, volume 12170 of *Lecture Notes in Computer Science*, pages 339–368, Santa Barbara, CA, USA, August 17–21, 2020. Springer, Cham, Switzerland.

[BCPZ16] Alberto Battistello, Jean-Sébastien Coron, Emmanuel Prouff, and Rina Zeitoun. Horizontal side-channel attacks and countermeasures on the ISW masking scheme. In Benedikt Gierlichs and Axel Y. Poschmann, editors, *Cryptographic Hardware and Embedded Systems – CHES 2016*, volume 9813 of *Lecture Notes in Computer Science*, pages 23–39, Santa Barbara, CA, USA, August 17–19, 2016. Springer, Berlin, Heidelberg, Germany.

[BDF+17] Gilles Barthe, François Dupressoir, Sebastian Faust, Benjamin Grégoire, François-Xavier Standaert, and Pierre-Yves Strub. Parallel implementations of masking schemes and the bounded moment leakage model. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *Advances in Cryptology – EUROCRYPT 2017, Part I*, volume 10210 of *Lecture Notes in Computer Science*, pages 535–566, Paris, France, April 30 – May 4, 2017. Springer, Cham, Switzerland.

[BDF24] Gianluca Brian, Stefan Dziembowski, and Sebastian Faust. From random probing to noisy leakages without field-size dependence. In Marc Joye and Gregor Leander, editors, *Advances in Cryptology – EUROCRYPT 2024, Part IV*, volume 14654 of *Lecture Notes in Computer Science*, pages 345–374, Zurich, Switzerland, May 26–30, 2024. Springer, Cham, Switzerland.

[BFO+21] Gianluca Brian, Antonio Faonio, Maciej Obremski, João L. Ribeiro, Mark Simkin, Maciej Skórski, and Daniele Venturi. The mother of all leakages: How to simulate noisy leakages via bounded leakage (almost) for free. In Canteaut and Standaert [CS21], pages 408–437.

[BFO23] Francesco Berti, Sebastian Faust, and Maximilian Orlt. Provable secure parallel gadgets. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2023(4):420–459, 2023.

[BMRT22] Sonia Belaïd, Darius Mercadier, Matthieu Rivain, and Abdul Rahman Taleb. IronMask: Versatile verification of masking security. In *2022 IEEE Symposium on Security and Privacy*, pages 142–160, San Francisco, CA, USA, May 22–26, 2022. IEEE Computer Society Press.

[BRT21] Sonia Belaïd, Matthieu Rivain, and Abdul Rahman Taleb. On the power of expansion: More efficient constructions in the random probing model. In Canteaut and Standaert [CS21], pages 313–343.

[BRTV21] Sonia Belaïd, Matthieu Rivain, Abdul Rahman Taleb, and Damien Vergnaud. Dynamic random probing expansion with quasi linear asymptotic complexity. In Mehdi Tibouchi and Huaxiong Wang, editors, *Advances in Cryptology – ASIACRYPT 2021, Part II*, volume 13091 of *Lecture Notes in Computer Science*, pages 157–188, Singapore, December 6–10, 2021. Springer, Cham, Switzerland.

[BS21] Olivier Bronchain and François-Xavier Standaert. Breaking masked implementations with many shares on 32-bit software platforms. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2021(3):202–234, 2021.

[CFOS21] Gaëtan Cassiers, Sebastian Faust, Maximilian Orlt, and François-Xavier Standaert. Towards tight random probing security. In Tal Malkin and Chris Peikert, editors, *Advances in Cryptology – CRYPTO 2021, Part III*, volume 12827 of *Lecture Notes in Computer Science*, pages 185–214, Virtual Event, August 16–20, 2021. Springer, Cham, Switzerland.

[CJRR99] Suresh Chari, Charanjit S. Jutla, Josyula R. Rao, and Pankaj Rohatgi. Towards sound approaches to counteract power-analysis attacks. In Michael J. Wiener, editor, *Advances in Cryptology – CRYPTO'99*, volume 1666 of *Lecture Notes in Computer Science*, pages 398–412, Santa Barbara, CA, USA, August 15–19, 1999. Springer, Berlin, Heidelberg, Germany.

[Cor14] Jean-Sébastien Coron. Higher order masking of look-up tables. In Nguyen and Oswald [NO14], pages 441–458.

[CPRR14] Jean-Sébastien Coron, Emmanuel Prouff, Matthieu Rivain, and Thomas Roche. Higher-order side channel security and mask refreshing. In Shiho Moriai, editor, *Fast Software Encryption – FSE 2013*, volume 8424 of *Lecture Notes in Computer Science*, pages 410–424, Singapore, March 11–13, 2014. Springer, Berlin, Heidelberg, Germany.

[CS21] Anne Canteaut and François-Xavier Standaert, editors. *Advances in Cryptology – EUROCRYPT 2021, Part II*, volume 12697 of *Lecture Notes in Computer Science*, Zagreb, Croatia, October 17–21, 2021. Springer, Cham, Switzerland.

[DDF14] Alexandre Duc, Stefan Dziembowski, and Sebastian Faust. Unifying leakage models: From probing attacks to noisy leakage. In Nguyen and Oswald [NO14], pages 423–440.

[DFS15a] Alexandre Duc, Sebastian Faust, and François-Xavier Standaert. Making masking security proofs concrete - or how to evaluate the security of any leaking device. In Elisabeth Oswald and Marc Fischlin, editors, *Advances in Cryptology – EUROCRYPT 2015, Part I*, volume 9056 of *Lecture Notes in*

Computer Science, pages 401–429, Sofia, Bulgaria, April 26–30, 2015. Springer, Berlin, Heidelberg, Germany.

[DFS15b] Stefan Dziembowski, Sebastian Faust, and Maciej Skorski. Noisy leakage revisited. In Elisabeth Oswald and Marc Fischlin, editors, *Advances in Cryptology – EUROCRYPT 2015, Part II*, volume 9057 of *Lecture Notes in Computer Science*, pages 159–188, Sofia, Bulgaria, April 26–30, 2015. Springer, Berlin, Heidelberg, Germany.

[DFZ19] Stefan Dziembowski, Sebastian Faust, and Karol Zebrowski. Simple refreshing in the noisy leakage model. In Steven D. Galbraith and Shiho Moriai, editors, *Advances in Cryptology – ASIACRYPT 2019, Part III*, volume 11923 of *Lecture Notes in Computer Science*, pages 315–344, Kobe, Japan, December 8–12, 2019. Springer, Cham, Switzerland.

[GP99] Louis Goubin and Jacques Patarin. DES and differential power analysis (the "duplication" method). In Çetin Kaya Koç and Christof Paar, editors, *Cryptographic Hardware and Embedded Systems – CHES'99*, volume 1717 of *Lecture Notes in Computer Science*, pages 158–172, Worcester, Massachusetts, USA, August 12–13, 1999. Springer, Berlin, Heidelberg, Germany.

[ISW03] Yuval Ishai, Amit Sahai, and David Wagner. Private circuits: Securing hardware against probing attacks. In Dan Boneh, editor, *Advances in Cryptology – CRYPTO 2003*, volume 2729 of *Lecture Notes in Computer Science*, pages 463–481, Santa Barbara, CA, USA, August 17–21, 2003. Springer, Berlin, Heidelberg, Germany.

[IUH22] Akira Ito, Rei Ueno, and Naofumi Homma. On the success rate of side-channel attacks on masked implementations: Information-theoretical bounds and their practical usage. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *ACM CCS 2022: 29th Conference on Computer and Communications Security*, pages 1521–1535, Los Angeles, CA, USA, November 7–11, 2022. ACM Press.

[JMB24] Vahid Jahandideh, Bart Mennink, and Lejla Batina. An algebraic approach for evaluating random probing security with application to AES. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2024(4):657–689, 2024.

[Man23] Giuseppe Manzoni. Reframing and extending the random probing expandibility to make probing-secure compilers tolerate a constant noise. Cryptology ePrint Archive, Paper 2023/1373, 2023.

[MMMS23] Loïc Masure, Pierrick Méaux, Thorben Moos, and François-Xavier Standaert. Effective and efficient masking with low noise using small-mersenne-prime ciphers. In Carmit Hazay and Martijn Stam, editors, *Advances in Cryptology – EUROCRYPT 2023, Part IV*, volume 14007 of *Lecture Notes in Computer Science*, pages 596–627, Lyon, France, April 23–27, 2023. Springer, Cham, Switzerland.

[MRS22] Loïc Masure, Olivier Rioul, and François-Xavier Standaert. A nearly tight proof of duc et al.'s conjectured security bound for masked implementations. In Ileana Buhan and Tobias Schneider, editors, *Smart Card Research and Advanced Applications - 21st International Conference, CARDIS 2022, Birmingham, UK, November 7-9, 2022, Revised Selected Papers*, volume 13820 of *Lecture Notes in Computer Science*, pages 69–81. Springer, 2022.

[MS23] Loïc Masure and François-Xavier Standaert. Prouff and Rivain's formal security proof of masking, revisited - tight bounds in the noisy leakage model. In Helena Handschuh and Anna Lysyanskaya, editors, *Advances in Cryptology – CRYPTO 2023, Part III*, volume 14083 of *Lecture Notes in Computer Science*, pages 343–376, Santa Barbara, CA, USA, August 20–24, 2023. Springer, Cham, Switzerland.

[NO14] Phong Q. Nguyen and Elisabeth Oswald, editors. *Advances in Cryptology – EUROCRYPT 2014*, volume 8441 of *Lecture Notes in Computer Science*, Copenhagen, Denmark, May 11–15, 2014. Springer, Berlin, Heidelberg, Germany.

[PGMP19] Thomas Prest, Dahmun Goudarzi, Ange Martinelli, and Alain Passelègue. Unifying leakage models on a Rényi day. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology – CRYPTO 2019, Part I*, volume 11692 of *Lecture Notes in Computer Science*, pages 683–712, Santa Barbara, CA, USA, August 18–22, 2019. Springer, Cham, Switzerland.

[PR13] Emmanuel Prouff and Matthieu Rivain. Masking against side-channel attacks: A formal security proof. In Thomas Johansson and Phong Q. Nguyen, editors, *Advances in Cryptology – EUROCRYPT 2013*, volume 7881 of *Lecture Notes in Computer Science*, pages 142–159, Athens, Greece, May 26–30, 2013. Springer, Berlin, Heidelberg, Germany.

[RP10] Matthieu Rivain and Emmanuel Prouff. Provably secure higher-order masking of AES. In Stefan Mangard and François-Xavier Standaert, editors, *Cryptographic Hardware and Embedded Systems – CHES 2010*, volume 6225 of *Lecture Notes in Computer Science*, pages 413–427, Santa Barbara, CA, USA, August 17–20, 2010. Springer, Berlin, Heidelberg, Germany.

[WJSW24] Bohan Wang, Fanjie Ji, Yiteng Sun, and Weijia Wang. Random probing security with precomputation. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2025(1):523–551, Dec. 2024.